

Safe Protocols for Generating Power Pulses with Heterogeneous Populations of Thermostatically Controlled Loads

Nikolai A. Sinitsyn

Los Alamos National Laboratory, Los Alamos, NM 87545, USA

Soumya Kundu*

University of Michigan, Ann Arbor, MI 48105, USA

Scott Backhaus

Los Alamos National Laboratory, Los Alamos, NM 87545, USA

Abstract

We explore methods to use thermostatically controlled loads (TCLs), such as water heaters and air conditioners, to provide ancillary services by assisting in balancing generation and load. We show that by adding simple imbedded instructions and a small amount of memory to temperature controllers of TCLs, it is possible to design open-loop control algorithms capable of creating short-term pulses of demand response without unwanted power oscillations associated with temporary synchronization of the TCL dynamics. By moving a small amount of intelligence to each of the end point TCL devices, we are able to leverage our knowledge of the time dynamics of TCLs to shape the demand response pulses for different power

*Corresponding author

Email addresses: nsinitsyn@lanl.gov (Nikolai A. Sinitsyn),
soumyak@umich.edu (Soumya Kundu), backhaus@lanl.gov (Scott Backhaus)

system applications. A significant benefit of our open-loop method is the reduction from two-way to one-way broadcast communication which also eliminates many basic consumer privacy issues. In this work, we focus on developing the algorithms to generate a set of fundamental pulse shapes that can subsequently be used to create demand response with arbitrary profiles. Demand response control methods, such as the one developed here, open the door to fast, nonperturbative control of large aggregations of TCLs.

Keywords: Thermostatically-controlled-loads, demand side control, ancillary services, generation-load balancing, hysteresis-based control

1. Introduction

Conventional thermal generators, which are constrained by physical ramp rates, have difficulty in manoeuvering to compensate for variability in loads or other generation, e.g. variability due to time-intermittent generation such as wind generation. On the other hand, electrical loads, e.g. thermostatically-controlled-loads (TCLs), may be a better alternative to provide the required generation-balancing ancillary services [1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15] because it is feasible for these loads to compensate for power imbalances much more quickly than thermal generators. TCLs such as air conditioners (AC) and water heaters account for about 50% of electricity consumption in the United States [16] implying that the potential resource provided by TCLs is quite large. A residential AC unit consumes 0.5-7 kW, and when aggregated over a large city with a million houses, the total controllable resource approaches ~ 3 GW. The typical TCL duty cycle will limit the amount of resource available at any one time, but even a relatively low duty cycle of 10% still yields a controllable resource of ~ 300 MW.

Because these TCL loads are on/off loads and not subject to the ramping limits of conventional thermal generation, they can be fully deployed on the time scale of the latency of the communication system used to access them, e.g. ~ 0.3 minutes [3, 17, 18, 19, 20]. Such short time scales for full deployment are not achievable with spinning or non-spinning reserves and, therefore, TCLs have considerable potential to provide both system reserves and generation-load balancing ancillary services [21, 22].

Past research on controlling TCLs has been either on direct interruption for long (several hours) duration [18, 23, 24, 25] or on shifting of the TCLs temperature set point [21, 26, 27, 28]. The infrequent application of the former technique is useful for critical peak shaving, but frequent application would quickly lead to customer fatigue. An approach based on small shifts of the TCL set point appears more attractive. In this case, subtle ($0.1\text{-}0.5^\circ\text{C}$) changes in the thermostat set point temperature can create significant changes in aggregate power consumption. Such small set point shifts would be almost unnoticed by customers and could be applied frequently as long there is not a long-term drift in one direction. However, this method is not without its difficulties. Extensive analytical and numerical modeling studies have shown that suddenly imposed changes to the operation cycles of a large population of TCL loads leads to temporary synchronization of otherwise uncorrelated autonomous devices resulting in potentially detrimental power oscillations [21, 26, 27, 28, 29]. Efforts to resolve this problem by feedback control techniques encounter additional problems, e.g. the need for two-way communication to provide full state information for each TCL and subsequent individual control of all TCLs. This approach, at least, requires development of new software that would process large amount of data on-fly while the state estimation must be

done in a noisy environment created by all the other loads on the system [30]. Perhaps more important is that two-way communication leads to privacy security issues that have already raised a popular protest [31, 32, 33, 34]. It seems that it can be implemented only on the voluntary basis, possibly with only a small fraction of TCL owners willing to lower their security level in return for rebate. For example, if only 5% from 1 million TCL owners in the area allow for the two-way communication with their utility company, this would lead to insignificant available power for control. However, the information collected from this 5% of the population should be sufficient to estimate some important aggregated parameters of the full population, such as the fraction of the currently working TCLs, the total power consumed by these TCLs, distribution of the temperature band widths of their operation, fraction of TCLs in the ON and OFF modes e.t.c. Extraction of this information, arriving to the utility company from only 5% of all TCLs, does not require a complex software and can be achieved "on-fly". Therefore, we believe that the technology that would allow for reliable one-way communication control of the remaining "95% of TCLs", and that would be based on the knowledge of only basic aggregate parameters of the full population of TCLs, should be attractive for large scale commercial application.


In this article, we propose an open loop approach which eliminates the problem of unwanted synchronization of TCLs without a need for two-way communication opening a path for implementing control functions with a heterogeneous population of TCLs using existing technologies. We show that by adding a small amount of intelligence (in the form of a few new elementary instructions) to the endpoint TCLs, we are able to implement simple open-loop controls which do not result in subsequent TCL load synchronization. Such strategies to control large

heterogeneous TCL populations, which we will call the safe protocols (SPs), could be quickly implemented to compensate for a generation-load mismatch without the fear of potentially dangerous subsequent TCL behavior. In [35], one such a SP for shifting temperature set point of the population of loads, without initiating any unwanted power oscillations, was proposed. The mechanism was shown to provide effective control on a relatively slow time scale (e.g. \sim 30 minutes). Here we extend the idea and show that embedded instructions allow to “safely” generate a class of power pulses of several basic shapes. For example, we will show that it is possible to generate a sharp (e.g. 3 minutes long) power pulse with subsequent slow return of all TCLs to the initial customer set values of the temperature band. At the other extreme, we demonstrate another safe protocol that generates time-extended pulses (\sim 0.5-1 hour) with changes in the TCL controlled temperatures of below 1°C. Such pulses are useful in spinning reserve applications. In this work, we focus on the creation of these fundamental shapes and we postpone the discussion of optimal pulse reconstruction and state estimation for future studies.

Our article is arranged as follows: In Section 2, we will outline the major problems that appear with straightforward (not safe) strategies to control the power of TCLs; in Section 3, a demonstration of a SP is presented. In Section 4, we extend the idea and present a second protocol that achieves a shift of temperature set points of TCLs; Section 5 presents a SP that creates a sharp power pulse of short duration; while Section 6 briefly reviews possible further strategies to combine different SPs to achieve a desired power response. We summarize our results in Section 7.

2. Problem Description

Before we proceed with introduction to SPs, we outline the basic problems encountered when more straightforward “unsafe” approaches are applied to control the power consumption of heterogeneous TCL populations.



(a) Dynamics of temperature (*top*) and consumed power (*bottom*) of a thermostatic load (b) TCL probability distribution function (PDF)

Figure 1: (left) Individual dynamics and (right) aggregate steady state probability distribution of TCLs over temperature in the OFF (red) and ON (blue) modes. Small nonzero density of TCLs beyond the deadband limits in (b) is due to the combined effect of temperature fluctuations and a finite response time to temperature measurements.

When not subject to major thermal disturbances, a TCL’s power consumption is cyclic in nature. If we consider a TCL providing cooling, the controlled temperature θ drops when the TCL draws power (in the ON state) and increases when the TCL is not drawing power (in the OFF state). Evolution of temperature in a

single house is generally modeled by a differential equation [26]:

$$\dot{\theta} = \begin{cases} -\frac{1}{CR}(\theta - \theta_{amb} + PR), & \text{ON state} \\ -\frac{1}{CR}(\theta - \theta_{amb}), & \text{OFF state} \end{cases} \quad (1)$$

where θ_{amb} is the ambient temperature, C is the thermal capacitance, R is the

- Current OFF state PDF
- Current ON state PDF
- - - Previous OFF state PDF
- - - Previous ON state PDF
- Direction of migration of OFF TCLs
- ← Direction of migration of ON TCLs
- ↗ Natural ON-to-OFF switching at lower boundary
- ↘ Natural OFF-to-ON switching at upper boundary
- ↑ Forced ON-to-OFF switching
- ↓ Forced OFF-to-ON switching

Figure 2: A list of different signs used in the following figures.

thermal resistance, and P is the cooling power of the TCL when in the ON state. P is proportional to the power drawn from the Power Grid with a coefficient that describes the efficiency of the TCL. We assume the regime of TCL operation with a duty cycle less than 100%. Such a TCL switches its state from OFF to ON when its temperature increases to θ_+ and from OFF to ON when temperature drops to θ_- . Such dynamics and control push the TCL into a hysteresis around the set point temperature $\theta_s = (\theta_- + \theta_+)/2$ with a dead band of $\Delta = (\theta_+ - \theta_-)$. Solution of Eqs. (1) with such boundary conditions is shown in Fig. 1(a) (black curve). Fig. 1(a) (red and blue curves) shows that at steady state the power consumption of a TCL switches between a constant value and zero in ON and OFF states, respectively. The heterogeneity and independence of the TCL dynamics ensure

that the time phase of these oscillations are not correlated and do not result in noticeable aggregate power oscillations. If we start from a set of arbitrary TCL initial states, the heterogeneity of C , R and P values across the population and natural temperature fluctuations will cause the aggregate power demand of a large ensemble of TCLs settle down to a constant value with fluctuations due to the finite size of the population. In this steady state, temperatures of all the TCLs would lie within their individual hysteresis dead bands and the temperatures of the TCLs attain steady state probability distribution [21], as presented in Fig. 1(b). For simplicity, we will use this steady state as our reference state for demonstration of response of the population to control signals. To help understand the figures to follow, a list of the different marks/signs used is presented in Fig. 2.

While undisturbed, the TCL dynamics remain uncorrelated, however, some forms of control can result in potentially dangerous correlations. For example, suppose that our goal is to reduce power consumed by all TCLs during a ten minute period and then return all TCLs to their individual dead bands. Naively, one may try switching off all TCLs for 10 minutes and then sending a signal to all of them to start operating according to the initial conditions [17, 18, 3, 19]. Fig. 3 shows that such a protocol would create a non-steady-state distribution with “kinks” and “gaps” in the ON-and OFF-state distributions.

In Fig. 4(a) we show our simulations of behavior of a population of 10000 TCLs¹ in response to such a control signal. At $t = 15$ hours, all the ON-state

¹We set R , P and C to follow the lognormal distribution in the ensemble of TCLs, which was argued to be a reasonable approximation to a real situation [21, 27], and we chose them to achieve a typical TCL cycle duration of approximately 45 minutes with a typical duty cycle approximately 40% at $\theta_{amb} = 32^{\circ}\text{C}$ and $\theta_s = 20^{\circ}\text{C}$. Number of TCLs is $N = 10,000$; width




Figure 3: Evolution of the TCL population distribution (y-axis) over temperature (x-axis) when all TCLs are forced to stay OFF for a certain duration: (a) On receiving signal, all TCLs switch to OFF state. (b) The population distribution simply keeps moving to the right, crossing the upper deadband limit, with no natural switching at either of the deadband limits. (c) At the end of the desired duration, a signal is sent to TCLs to return to their original state which results, in particular, in all TCLs with $\theta > \theta_+$ turning ON. (d) After all TCLs return inside the originally set temperature range, a kink (higher than normal density of TCLs) and a hole (zero density over temperature) circulate with time. Notation is explained in Fig. 2




Figure 4: (a)(left) Aggregate power consumption when all TCLs are forced to switch OFF for 10 min at 15 hrs. (a)(right) Red and blue curves show temperature evolution of two individual TCLs that are OFF and ON, respectively, at 15 hrs. (b) Same as in (a) but for the case when initially all TCLs are requested to switch ON for 10 min.

TCLs are switched OFF and all the TCLs are kept in switched OFF state for 10 min. The aggregate TCL power, in Fig. 4(a)(left), drops to zero, however, after receiving the signal to return to the initial state, the aggregate power does not return to the initial steady state value (60MW). Instead, the population consumes much more power (with a peak above 100MW). During the ten minutes without the temperature control, many of the controlled temperatures increased beyond the customer upper dead band θ_+ , as shown in Fig. 4(a)(right), and a sudden signal to return to the initial temperature range forces a larger fraction of TCLs to switch

of the deadband $\Delta = 1 \text{ } ^\circ\text{C}$; Parameters C, R and P have mean values $3 \text{ } \text{kW}/\text{ } ^\circ\text{C}$, $2 \text{ } \text{ } ^\circ\text{C}/\text{kW}$ and $14 \text{ } \text{kW}$, respectively, and a standard deviation of $\sigma_p = 0.07$ of the corresponding mean; a zero-mean Gaussian noise with standard deviation of $0.052 \text{ } ^\circ\text{C}/\text{min}^{0.5}$ was added to (1) to account for natural temperature fluctuations.

to the ON state than had originally switched to the OFF state. This phenomenon is known as a *cold load pickup* [36, 37]. Similar synchronization happens when TCLs are switched ON for a 10-min interval (Fig. 4(b)).

After providing beneficial control during the first ten minutes, this naive scheme creates an uncontrolled power spike of a similar magnitude and opposite sign. Perhaps even worse, the following power consumption relaxes to the steady state not monotonously but rather showing pronounced power oscillations, which may persist for several cycles depending on the population’s heterogeneity. These oscillations, which we will call *parasitic oscillations*, are sustained by the synchronization of individual phases of TCLs induced by simultaneous transition of a large number of TCLs into the same ON-state. Fig. 4 (left) shows the inverse situation when all the TCLs are set to be ON for ten minutes after which the originally-ON-TCLs are switched back from OFF to ON resulting in a similar damped oscillation.

Another “naive” approach to control the aggregated power of TCLs is to impose a sudden shift of dead band positions [21, 35], as shown in Fig. 5. This approach of a sudden shift in deadband limits uniformly across the population also creates kinks and gaps in the temperature distribution, which sustain for a long duration. Fig. 6(a) shows our numerically calculated aggregated power response of a group of TCLs, which were initially operating in a steady state, after their temperature set point was increased by $0.4\text{ }^{\circ}\text{C}$. As a result, some of the TCLs instantly found themselves lying below θ_- , as shown in Fig. 6(b), and therefore switched OFF resulting in an immediate drop in the aggregate demand. After this immediate drop, however, the aggregate demand undergoes large damped oscillatory mode due to essentially the same reasons as in previous example.




Figure 5: Evolution of temperature distribution under sudden shift of deadband, by δ , from $[\theta_-^0, \theta_+^0]$ to $[\theta_-, \theta_+]$: (a) As the deadband is shifted, a fraction of the initially OFF TCLs "instantly" switches to ON, creating a "kink" in the OFF-distribution. (b) While OFF-state TCLs are yet to reach the new upper deadband limit, ON-to-OFF switching already occurs at the new lower deadband limit. (c) As switching starts at both the new deadband limits, a "hole" is created in the ON-distribution. (d) The "kink" and "hole" keep moving around the distribution.




Figure 6: (a) Aggregate response (left) under temperature setpoint shift (right) by $0.4\text{ }^{\circ}\text{C}$. (b) Individual temperature evolution of two TCLs.

In summary, the synchronization of TCLs by external control signals may induce considerable power oscillations with amplitudes comparable to the total power consumed by the population at steady conditions. Control strategies that do not avoid such unwanted parasitic power oscillations may be more harmful than the power fluctuations that they attempt to remove.

3. Safe protocol-1

We start our discussion of safe strategies to create power pulses in the grid by considering a practically less interesting but more easy-to-follow example. Consider the strategy to control the population, which is illustrated in Fig. 7 and which we will term the safe protocol-1, or simply SP-1. To implement any SP, an important assumption is that temperature controllers have additional embedded instructions. Specifically for SP-1, first, we assume that after receiving the control signal, TCLs are capable to save the information about the temperature that they

are having at this moment. Second, we assume that TCLs are capable to compare their current temperature with their saved one. Finally, we assume that by reaching such a saved temperature, the TCLs can make transitions between ON and OFF states purely autonomously, i.e. without waiting for an external signal. Such additional instructions are elementary and can be easily added to temperature controllers of TCLs in most houses.




Figure 7: Distribution of TCLs (y-axis) over temperature (x-axis) in the ON and OFF modes during stages of safe protocol-1 (SP-1). (a) All TCLs are requested to switch OFF. (b) and (c) Intermediate stages of evolution of population distribution (y-axis is not to scale). (d) After all TCLs appear again in the OFF mode, TCLs that were in the ON mode before application of the protocol make autonomous transition from OFF to ON. As a result, the original steady state distribution is reproduced.

Having such additional capabilities, the generation of a safe pulse can proceed as follows:

1. By receiving an external signal, TCLs from the ON branch switch to OFF (Fig. 7(a)).
2. At this moment, TCLs that switched from ON to OFF save their temperatures.
3. All TCLs of the ensemble continue working without changing the position of the temperature operation band, as shown in Fig. 7(b)-(c).
4. When temperature of TCLs that switched OFF at step-1 reaches the same temperature as the saved one when it is again in the OFF-mode (this stage is achieved after one period of the cycle, as shown in Fig. 7(d)), these TCLs autonomously switch to the ON state.

A similar process works if the power needs to be increased at initial moment, by switching TCL's states, initially, from OFF to ON at step-1.

Fig. 8 shows two simulation results of application of SP-1 to instantaneously decrease (Fig. 8(a)) and increase (Fig. 8(b)) power in a heterogeneous population of TCLs. While after the immediate change, the aggregate power recovers by making a swing in the opposite direction, it quickly settles down back to its original steady state value after one cycle duration of a typical TCL in the population. This becomes possible because the original distribution is reproduced after completion of stages of SP-1 (Fig. 7(d)). Fig. 8(a)(right) illustrates that, while TCL_1 that was OFF at the start of SP-1 at time 15 hrs continues with its natural cycle, TCL_2 switched its state from ON to OFF and then switched back from OFF to ON after a duration equal to its natural cycle period. An equivalent behavior occurs when TCLs are initially switched ON, Fig. 8(b).

Here we note that Fig. 7 shows that the SP-1 ends with equilibrium distribution of TCLs over temperature if the population is homogeneous, i.e. if all parameters




Figure 8: Aggregate power (left) and individual temperature (right) response of TCLs under SP-1. Notation is explained in caption for Fig. 4.

of TCLs are identical with identical cycle duration. A real population would be certainly heterogeneous. Nevertheless, Fig. 8 for the power response shows that SP-1 does not lead to unwanted power oscillations by the end of the protocol even if it is applied to a *large* heterogeneous population described in the footnote in Section II. The reason is that for any TCL in a large heterogeneous population, there are many other TCLs with similar parameters. That is, each large population can be considered as a set of many almost homogeneous sub-populations. Since a SP does not lead to unwanted power oscillations in *any* homogeneous sub-population, it does not induce such oscillations when it is applied to a full heterogeneous ensemble of TCLs either.

Safe protocol SP-1 has two additional useful features:

- Power increases almost instantly in response to the external signal so that the protocol can be used to offset fast power fluctuations. In reality, of course, there is always a time delay between receiving the signal and chang-

ing the mode of operation of a TCL, which restricts the speed of control. Response time of a realistic TCL to external control, however, is as small as a fraction of a minute [19].

- Since the temperature deadband remains undisturbed, this protocol guarantees that we never leave the temperature operation band that was set by a customer. Hence, application of this protocol should be practically unnoticeable by owners of TCLs.
- Finally, and this will be a generic property of all safe protocols discussed in this article, the safe protocols are also safe from the privacy point of view. They do not require the utility company to have information about any specific TCL. Only aggregate information about the total power consumed by the ensemble is needed to estimate the power response to a SP.

However, the applicability of SP-1 alone is limited because SP-1 cannot generate any extra energy by the time of its completion in comparison to the energy generated by the same ensemble at the steady state during the same time. At the end of SP-1, the ensemble of TCLs is brought to the initial distribution, exactly the same state in which ensemble would be if no external control is applied. Intuitively, the conservation of energy argument suggests that SP-1, while producing a nontrivial power pulse, does not save any extra energy by the time of its completion, which is indeed the case as we prove in the Appendix.

4. Safe Protocol-2

The next protocol, SP-2, shows that the drawback of the previous protocol, i.e. the inability to store/release energy by the time of its completion, can be

eliminated by choosing a different strategy. In particular, SP-2 creates a power pulse with a net release of energy, in comparison to the undisturbed ensemble, to the grid. This is achieved during a shift of the positions of TCL's temperature operation bands by a fixed amount of temperature. This protocol has been already presented previously [35], so we only describe it here briefly because it will be used later in our discussion.

To describe the SP-2, we should introduce the *transition points* that are *intermediate* hysteresis deadband limits that become operative starting from the instant the temperature shift is initiated until the new set of deadband limits are applied across the population. Let θ_-^0 and θ_+^0 be the lower and upper deadband limits before a shift of the deadband position to the new limits $\theta_- = (\theta_-^0 + \delta)$ and $\theta_+ = (\theta_+^0 + \delta)$. If the temperature setpoint is to be increased by an amount δ , then the transition points would be θ_-^0 and $(\theta_+^0 + \delta)$. While for a decrease in setpoint by an amount δ , the transition points would be $(\theta_-^0 - \delta)$ and θ_+^0 . Note that positions of the transition temperature points do not coincide with the positions of the new dead band limits.

The SP-2 was constructed as follows:

1. A TCL continues to stay in its present state (ON or OFF) until it hits one of the transition points.
2. Once it has reached one of the transition points, the new pair of deadband limits, $\theta_- = (\theta_-^0 + \delta)$ and $\theta_+ = (\theta_+^0 + \delta)$, starts governing its thermal dynamics.

Here $\delta > 0$ for the increase in the setpoint and $\delta < 0$ for the decrease in the setpoint. Fig. 9 shows how this protocol works when the setpoint temperature is

increased by δ . The aggregate power demand of the population of TCLs attains a new steady state corresponding to the shifted deadband positions within a thermostat's cycle duration. Fig. 10 illustrates how SP-2 performs when the setpoint is increased. The aggregate response in Fig. 10(a) shows that the power drops for a duration equal to the time period and then stabilizes to the new steady state value without any parasitic oscillation. Fig. 10(b) shows two typical responses of individual TCLs. TCL₂ was initially OFF and hence simply continues with the new deadband limits, while TCL₁ which was ON at the initiation of SP-2 continues with its old deadband limits until it hits the lower deadband limit after which it adopts the new limits.




Figure 9: Evolution of TCL's distribution over temperature in the ON and OFF modes during states of Safe protocol-2 (SP-2), which is based on shifting temperature range limits from initial, (θ_-^0, θ_+^0) , to new ones, (θ_-, θ_+) .




Figure 10: Aggregate power and individual temperature response of TCLs under SP-2 when the temperature setpoint is increased by $0.4\text{ }^{\circ}\text{C}$. Notation is explained in caption for Fig. 6.

Besides the absence of parasitic oscillations, this safe protocol (SP-2) has several new advantages:

- It can be used to switch the set point temperatures almost in arbitrary range, limited only by customer tolerance level.
- This protocol practically does not require extra switching of the device branch per time of the driving protocol. The completion of the protocol takes about the cycle duration time and by this time all TCLs switch their mode (from OFF to ON or from ON to OFF) only one time. In comparison, SP-1 takes about the same time before its completion during which each TCL makes 4 mode switching events, while a steady state cycle includes only 2 switching events between OFF and ON states per one TCL cycle. This is important because more frequent switching of the TCL operation mode accelerates the fatigue of the device circuits and may be considered the major cost of application of the control.

By careful examination of Fig. 9, one can recognize the reason why SP-2 does not lead to unwanted power oscillation: All steps are chosen to avoid creation of "kinks" or "holes" in the final distribution of TCLs over temperature.

The major drawback of SP-2 is its "slowness". Power does not change abruptly after receiving a control command. It takes a certain while to reach the extremum of power consumption by an ensemble. Hence SP-2, alone, is restricted to very specific applications that do not require fast response, e.g. for offsetting relatively slow (\sim 30-60 minutes in our case) but possibly large power fluctuations. For example, one can envision a situation when a power line breaks down and one wants to decrease the burden by TCLs on the Power Grid during a relatively large time interval, e.g. to provide time for a backup coal power plant to increase power output. Such situations are commonly resolved by relatively expensive, in comparison to standard power generators, spinning reserves, which annual cost on the scale of such a country as the USA is considerable. This justifies the importance of the SP-2 for helping to resolve such problems.

5. Safe Protocols to Generate Short Power Pulses

TCL's ability to respond to control signals at a fraction of a minute suggests that they may not be simply considered as replacement to standard spinning reserves. They can be also used to increase stability of the power supply in the power grid by offsetting sharp power spikes. For such applications, it is desirable to design control protocols that are safe against synchronization problems but react much faster and create very short power pulses. While safe protocols SP-1 and SP-2 could create power profiles without unwanted power oscillations, none of them is suitable for offsetting short (2-5 minutes) power spikes because the

pulses that they generate continue for a duration of a cycle of a typical TCL in the population.




Figure 11: SP-3 based on switching (to OFF) and shifting, followed by SP-2, to bring the population of TCL's to the initial setpoints.

Here we construct the protocol, referred to as SP-3, that generates a short power pulse of possibly large amplitude. If we consider that time of response of a power plant to fluctuations is e.g. 8 minutes, then what one would need to increase stability of the power supply is to be able to create shorter sharp pulses of duration, e.g., $\Delta t = 3$ minutes with a TCL ensemble and then return all TCLs to initial setpoints avoiding sharp pulses of opposite sign and power oscillations. This is achieved by the protocol, which stages are shown in Fig. 11. The mechanism of SP-3, when the need is to decrease power, is as follows:

1. By receiving the signal, TCLs in the ON mode switch to OFF (Fig. 11(a))

and save information about their original deadband limits, θ_-^0 and θ_+^0 .

2. All TCLs continue on OFF mode for a duration Δt (say, 3 min), as shown in Fig. 11(b). It is to be noted, that during this time Δt , no switching takes place at either of the deadband limits.
3. At the end of this interval, TCLs that were forced to switch the branch at step-1 now should switch from OFF to ON mode. The result is shown in Fig. 11(c).

Note that at this stage no “kinks” or “holes” in the distribution were created. Steps 1-3 will produce sharp drop of power consumption for a time Δt and resulting distribution at the end of step-3 will be close to the new equilibrium because, after step-3, we did not create any kinks or holes in TCL’s distribution. These steps, however, were also discussed in Section 2, in which we showed that if, following step 3, we simply ask all TCLs to return to initial conditions, parasitic power oscillations can be considerable. In order to fix this problem, we propose an application of SP-2 to effect a slow return of deadband limits.

4. After step-3, the ensemble is brought to the initial state (Fig. 11(d)) by application of a specific version of SP-2 (refer to Section 4). SP-2 that was described in the previous section requires information about the current deadband limits and the new deadband limits (or, alternatively the amount of shift in setpoint temperature), to use as transition points. But the “current” deadband temperature limits, at the end of step-3, θ_- and θ_+ are not exactly known for a heterogeneous population. One can show, however, that this knowledge is not needed. We can simply instruct all TCLs that switch to ON-branch after step-3 to start operation according to the old, saved, tem-

perature deadband (θ_-^0 to θ_+^0) after step-3. As for the TCLs that were in the OFF-branch before step-1, we instruct them to continue staying on this branch for a duration Δt from the moment they reach the temperature θ_+^0 (saved in step-1). This can be achieved by adding timers to the temperature controllers so that the timer starts autonomously after reaching θ_+^0 . Only after completing such extra Δt delay in the OFF-mode, these TCLs switch from OFF to ON and then set their thermostats to work according to the old band position. This set of instructions will not create kinks in the distribution and will be equivalent to application of SP-2, with effective transition point at θ_-^0 and θ_+ , following step-3. Note that step-4 is necessarily slow so that sharp and strong drop of power during steps 1-3 is compensated by a slow and weak increase of power during step-4.

Fig. 12 shows simulation results of the application of SP-3. A negative pulse of width 3 min is generated following the steps 1-4 of SP-3, as shown in Fig. 12 (right). The sharp and exact positive pulse is followed by a slow and weak change in aggregate power in the opposite direction for about a cycle duration before returning to the original steady state consumption level. A similar result, but in opposite direction, happens when a positive pulse of width 3 min is generated (Fig. 12 (left)).

This method is particularly suitable for applications, in which the required pulse width is much smaller than the cycle duration, e.g. less than 5 minutes. This strategy opens the possibility of an application of TCLs to a class of problems that standard spinning reserves cannot resolve.




Figure 12: Power response under SP-3 generating a positive pulse by switching ON (*left*) or negative pulse by switching OFF (*right*), both of a width of 3 min.

6. Hybrid protocols

In previous sections, we described three protocols that generate power pulses of different shapes with high degree of control. In real situations, those shapes separately may not be sufficient to achieve all utility goals - e.g. creating arbitrary shape of power response. The latter goal can be achieved by combining different pulses. For example, a train of power spikes can be eliminated by a train of SP-3 pulses sent to sub-populations of TCLs. However, applying SP-3 alone in a series of pulses may not be the best strategy for some goals. For example, if TCLs are used to replace standard spinning reserves, then the typical situation that they may need to resolve is to provide extra power supply at fixed level for relatively long duration, e.g. of 30 minutes. In this section, we show an example of a hybrid protocol that can be used in such a specific situation.

Suppose that in response to a damage of a power line, a spinning reserve should react almost instantly and supply extra power in the Power Grid during 30 minutes. In principle, extra power can be supplied by SP-2 but, alone, it is

slow. On the other hand, SP-1 responds "instantly" but does not supply an overall extra power. This suggests that if we split our ensemble of TCLs into two parts -




Figure 13: Combination of SP-1 and SP-2 into a hybrid protocol to generate a (negative) pulse. N is the total number of TCLs in the population, p is the fraction of TCLs that receive SP-1 signal and $(1 - p)$ is the fraction of TCLs that receive SP-2 signal.

one that receives SP-1 and one that receives SP-2 then we can generate a desired power profile that changes power consumption almost instantly and sustains it at the required level for sufficiently long duration. In Fig. 10 (left) and Fig. 8 (right) both power profiles of SP-1 and SP-2 take about a TCL cycle duration before settling down to a steady state value, and the power profiles experience peaks at around the same time in opposite directions. Thus if a section of a population fol-

lows SP-1, and the rest of TCLs follow SP-2, the aggregate power can be made to follow a square-like pulse profile before settling down after one cycle duration. Of course, the challenge is to design the population size undergoing protocols SP-1 and SP-2, and also the amount of setpoint shift to be used for SP-2. Fig. 13 shows a graphical illustration of how a hybrid SP is supposed to work.

The following parameters are important to quantify and design the protocol:

- p : the fraction of the population, of size N , undergoing SP-1 (while the rest of the population, of size $(N - Np)$ undergoes SP-2), and
- δ : the amount of the temperature shift to be used for SP-2. To generate a negative pulse, $\delta > 0$ (i.e. $\text{sign}(\delta) = 1$) and accordingly in SP-1 all the TCLs are to be switched OFF. Opposite actions are to be taken to generate positive pulse.

Fig. 14 presents two sample simulation results of application of this strategy for generating a positive (Fig. 14(a)) and a negative pulse (Fig. 14(b)). Fig. 14(b)




Figure 14: Power response under hybrid protocol that combines SP-1 and SP-2.

shows that by combining two SPs, we can instantly respond to emergency by decreasing power consumed by TCLs and sustain the lower power demand by the population almost at constant value during half an hour. During the following half an hour the power demand by the population is slowly growing and saturating close to the initial value. This is a type of the power profile that is frequently needed from spinning reserves [17, 18].

7. Conclusion

We have described and simulated the operation of a class of open-loop safe protocols that resolve the TCL synchronization problem nearly eliminating oscillations of controlled TCL power. These safe protocols have several benefits:

- absence of synchronization of TCLs and related parasitic oscillations
- fast control response (limited mostly by communication latency) enabling the generation of sharp pulses
- no requirement for two-way communication, and consequently preservation of privacy
- very limited information about the TCL ensemble is required to estimate the response.
- TCL ensembles can be for both up and down regulation, i.e. TCL power can be both increased and decreased

However, safe protocols also appear to have certain limitations whose impact require further investigation:

- more frequent than normal switchings of TCLs can shorten their lifetime. For example, when SP-3 is used with a small duration δt , a TCL will undergo at least two switchings while normally it takes a cycle duration (e.g. ~ 60 min in the studied population) to complete two switchings. This effect likely limits SP-3 to specific applications to responding to unusually large but relatively rare power imbalances such as on the shoulder of wind ramps
- the response of air conditioners to the control signal is usually restricted by a delay of about 20 seconds, which restricts duration and precision of SP-3 application to this class of TCLs
- customer tolerance can limit allowable changes in temperature setpoint restricting the application of SP-2
- the finite size of a TCL population restricts the total energy that can be compensated through pulse generation. After obtaining a useful power pulse, it takes about the natural TCL cycle time for the switched safe protocols to "recover", i.e. to bring switched TCLs back to the customer-defined dead band range. Therefore, a given TCL can be used only once per this characteristic time
- an estimate of at least some aggregated TCL ensemble parameters are still desired for accurate control of generated power response. This information is likely obtainable from monitoring at the utility level, e.g. by observing response from weak short pulses generated by SP-3. An alternative possibility is to use a *limited* number of houses with smart meters for two-way communication to provide information that is sufficient to estimate basic aggregated parameters of the whole population

- the safe protocols require a small amount of trivial to realize embedded intelligence in each TCL controller.

If these challenges can be overcome, the potential impact of TCL demand response, and specifically the safe protocols described here, would be large. Current estimates show that there are over 130 millions houses in the USA alone with over 100 millions of them equipped with air conditioners [38]. Considering a conservative estimate of 100MW reliable power control from air conditioners and a similar amount from water heaters in a one million houses, this technology implemented on the national scale could provide up to 20 GW power of ancillary services.

Acknowledgment

The authors would like to thank Professor Ian Hiskens of University of Michigan, Ann Arbor for his valuable guidance throughout this work. The work at LANL was carried out under the auspices of the National Nuclear Security Administration of the U.S. Department of Energy at Los Alamos National Laboratory under Contract No. DE-AC52-06NA25396.

Appendix A.

Theorem Appendix A.1. (No-Go Theorem) *One cannot generate extra energy by switching conditioners between ON and OFF branches along the SP-1 in a large heterogeneous ensemble of TCLs.*

Proof. Consider application of the SP-1 to a large population of TCLs, that are initially in the uncorrelated steady state. We assume that due to large size of the population, it can be partitioned by a number of sub-populations, which still have

large sizes, with practically identical parameters of TCLs and ambient conditions. It is then sufficient to prove the No-Go Theorem for an arbitrary homogeneous sub-population, because if any subpopulation does not produce an excess of energy then neither does their ensemble.

The evolution of temperature of a TCL, in (1), is linear both in temperature and consumed power. Hence, after averaging (1) over the population with random initial phases, we find that the averaged consumed from the Power Grid power $\langle P \rangle$, and the averaged temperature $\langle \theta \rangle$ of the ensemble also satisfy a linear relation:

$$\alpha \frac{d}{dt} \langle \theta \rangle = -\langle P(t) \rangle + \beta(\langle \theta \rangle - \gamma), \quad (\text{A.1})$$

where α , β , and γ are constants that characterize the ensemble. They can be expressed via the number of TLCs N , the efficiency of the TCLs, and parameters R , C and θ_{amb} . At the steady state, we would have $\frac{d}{dt} \langle \theta \rangle = 0$, so that the average power consumed by the ensemble in the steady state is $\langle P_{st} \rangle = \beta(\langle \theta_{st} \rangle - \gamma)$, where $\langle \theta_{st} \rangle$ is the average temperature in the unperturbed ensemble of TCLs. The total energy consumed by the ensemble in the steady state during some time τ then would be

$$W_{st}(\tau) = \beta(\langle \theta \rangle - \gamma)\tau. \quad (\text{A.2})$$

Consider now the total energy, W_{tot} , generated by the TCL ensemble during the time t_p between the initial and the final stages of the SP-1:

$$W_{\text{tot}} = \int_0^{t_p} dt \langle P(t) \rangle = \beta \int_0^{t_p} dt (\langle \theta \rangle - \gamma), \quad (\text{A.3})$$

where we used (A.1) and the fact from Section III that, for SP-1, at time t_p , the final distribution of TCLs over temperature and ON/OFF states coincides with the initial distribution. This means that the average temperature of the ensem-

ble coincides with the initial one, so that integrating the lhs of (A.1) we obtain

$$\int_0^{t_p} \frac{d}{dt} \langle \theta \rangle dt = 0.$$

Here we note that the time t_p coincides with the period of the cycle of a single TCL. Indeed, if we trace the evolution of any TCL of the ensemble during the SP-1 stages, we find that it goes purely autonomously during the period t_p through all stages of the considered temperature operation band, just like TCLs in nonperturbed ensemble except switching the branches at the beginning and the end of the protocol. We also note that at the steady state, the averaged over time power consumed by one TCL during the cycle of its operation coincides with the power averaged over the large population of TCLs with the random phase distribution. This means that, for a single TCL in a homogeneous population, $\int_0^{t_p} \theta dt = \langle \theta_{\text{st}} \rangle t_p$.

Hence, (A.3) can be rewritten as

$$W_{\text{tot}} = \beta(\langle \theta_{\text{st}} \rangle - \gamma)t_p, \quad (\text{A.4})$$

But (A.4) is exactly the energy W_{st} in (A.2) consumed by the TCL ensemble during time t_p without application of the external control. Hence the excess energy is $W_{\text{tot}} - W_{\text{st}} = 0$. \square

References

- [1] G. Strbac, Demand side management: Benefits and challenges, *Energy Policy* 36 (2008) 4419–4426.
- [2] M. Klobasa, Analysis of demand response and wind integration in Germanys electricity market, *IET Renewable Power Generation* 4 (2010) 55–63.
- [3] Heffner G., Goldman C., Kirby B. J., Kintner-Meyer M., Loads providing

ancillary services: review of international experience, Lawrence Berkeley National Laboratory, Berkeley (2007).

- [4] N. Lu, S. Katipamula, Control strategies of thermostatically controlled appliances in a competitive electricity market, IEEE Power Engineering Society General Meeting 1 (2005) 202–207.
- [5] K. Schisler, T. Sick, K. Brief, The role of demand response in ancillary services markets, Transmission and Distribution Conference and Exposition (2008) 1–3.
- [6] Z. Xu, J. Ostergaard, M. Togeby, Demand as frequency controlled reserve, IEEE Transactions on Power Systems 99 (2010) 1–10.
- [7] N. Ruiz, I. Cobelo, J. Oyarzabal, A direct load control model for virtual power plant management, IEEE Transactions on Power Systems 24 (2009) 959–966.
- [8] R. P. Malhamé, A statistical approach for modeling a class of power system loads, Ph. D. Thesis, Georgia Institute of Technology, USA, 1983.
- [9] A. Molina-García, M. Kessler, J. Fuentes, E. Gómez-Lázaro, Probabilistic characterization of thermostatically controlled loads to model the impact of demand response programs, IEEE Transactions on Power Systems 26 (2011) 241–251.
- [10] S. El-Ferik, S. Hussain, F. Al-Sunni, Identification of physically based models of residential air-conditioners for direct load control management, Proceedings of the 5th Asian Control Conference 3 (2004) 2079–2087.

- [11] W. Burke, D. Auslander, Robust control of residential demand response network with low bandwidth input, Proceedings of the ASME Dynamic Systems and Control Conference (2008) 1187–1189.
- [12] J. C. Laurent, G. Desaulniers, R. P. Malhamé, F. Soumis, A column generation method for optimal load management via control of electric water heaters, IEEE Transactions on Power Systems 10 (1995) 1389–1400.
- [13] R. Parsons, ASHRAE Handbook: Fundamentals, American Society of Heating, Refrigerating and Air-Conditioning Engineers (2005).
- [14] M. A. Zehir, M. Bagriyani, Demand Side Management by controlling refrigerators and its effects on consumers, Energy Conversion and Management 64 (2012) 238–244.
- [15] C. Ivarez Bel, M. A. Ortega, G. E. Escriv, A. G. Marn, Technical and economical tools to assess customer demand response in the commercial sector, Energy Conversion and Management 50 (2009) 2605–2612.
- [16] 2001 Residential Energy Consumption Survey, U.S. Energy Information Administration (2001).
- [17] B. Kirby, J. Kueck, T. Laughner, K. Morris, Spinning reserve from hotel load response, Electricity Journal 21 (2008).
- [18] B. Kirby, Load response fundamentally matches power system reliability requirements, IEEE Power Engineering Society General Meeting (2007).
- [19] B. Kirby, Spinning reserve from responsive loads, Technical Report ORNL/TM-2003/19, Oak Ridge National Laboratory, USA (2003).

- [20] Chellury Sastry, Viraj Srivastava, Rob Pratt, Shun Li, Use of Residential Smart Appliances for Peak-Load Shifting and Spinning Reserves: Cost/Benefit Analysis (2010).
- [21] D. S. Callaway, Tapping the energy storage potential in electric loads to deliver load following and regulation, with application to wind energy, *Energy Conversion and Management* 50 (2009) 1389–1400.
- [22] D. S. Callaway, I. A. Hiskens, Achieving Controllability of Electric Loads, *Proceedings of IEEE* 99 (2011) 184–199.
- [23] C. Y. Chong, A. S. Debs, Statistical synthesis of power system functional load models, *Proceedings of the 18th IEEE Conference on Decision and Control* 18 (1979) 264–269.
- [24] R. E. Mortensen, K. P. Haggerty, A stochastic computer model for heating and cooling loads, *IEEE Transactions on Power Systems* 3 (1988) 1213–1219.
- [25] R. Malhamé, C. Y. Chong, Electric load model synthesis by diffusion approximation of a high-order hybrid-state stochastic system, *IEEE Transactions on Automatic Control* 30 (1985) 854–860.
- [26] S. Kundu, N. Sinitsyn, S. Backhaus, I. A. Hiskens, Modeling and Control of Thermostatically-Controlled-Loads, *17th Power Systems Computations Conference* (2011).
- [27] C. Perfumo, E. Kofman, J. H. Braslavsky, J. K. Ward, Load management: Model-based control of aggregate power for populations of thermostatically controlled loads, *Energy Conversion and Management* 55 (2012) 36–48.

- [28] S. Bashash, H. K. Fathy, Modeling and Control Insights into Demand-side Energy Management through Setpoint Control of Thermostatic Loads, American Control Conference (2011) 4546 – 4553.
- [29] F. Koch, J. L. Mathieu, D. S. Callaway, Modeling and control of aggregated heterogeneous thermostatically controlled loads for ancillary services, 17th Power Systems Computations Conference (2011).
- [30] J. L. Mathieu, D. Callaway, State estimation and control of heterogeneous thermostatically controlled loads for load following, HICSS (2012) 2002–2011.
- [31] A. Damon, Smart meters spawn conspiracy talk: They know what you're watching on TV, Las Vegas SUN (2012).
- [32] Stop Smart Meters: The Film, (Online material: <http://www.stopsmartmetersnow.com/?p=132>) (2012).
- [33] Big Brother: The insidious Smart Meter, Radio Patriot (Online material: <http://radiopatriot.wordpress.com/2011/07/10/big-brother-the-insidious-smart-meter/>) (2011).
- [34] No More Privacy: Smart Meters Are Surveillance Devices That Monitor The Behavior In Your Home Every Single Minute Of Every Single Day, in: The American Dream (Online material: <http://endoftheamerican dream.com/archives/no-more-privacy-smart-meters-are-surveillance-devices-that-monitor-the-behavior-in-your-home-every-single-minute-of-every-single-day>).

- [35] S. Kundu, N. Sinitsyn, Safe Protocol for Controlling Power Consumption by a Heterogeneous Population of Loads, American Control Conference (2012).
- [36] S. Ihara, F. C. Schweppe, Physically based modelling of cold load pickup, IEEE Transactions on Power Apparatus and Systems 100 (1981) 4142–4150.
- [37] C. Y. Chong, R. P. Malhamé, Statistical synthesis of physically based load models with applications to cold load pickup, IEEE Transactions on Power Apparatus and Systems (1984) 1621–1628.
- [38] S. Cox, Cooling a Warming Planet: A Global Air Conditioning Surge, Yale Environment 360: Analysis (Online available at: e360.yale.edu/feature/cooling_a_warming_planet_a_global_air_conditioning_surge/2550) (2012).